

# **APPLICATION FOR UNITED STATES PATENT**

**in the name of**

**David Bill, Scott Brown, Gil Weigand, and Bill Raduchel**

**of**

**America Online**

**for**

**Generating Multiple Data Streams from a Single Data  
Source**

Fish & Richardson P.C.  
601 Thirteenth Street, NW  
Washington, DC 20005  
Tel.: (202) 783-5070  
Fax: (202) 783-2331

**ATTORNEY DOCKET:**

**06975-178001**

# Generating Multiple Data Streams from a Single Data Source

This application claims the benefit of U.S. Provisional Application No. 60/286,964 filed April 30, 2001, which is incorporated by reference.

## TECHNICAL FIELD

This invention relates to streaming media.

## BACKGROUND

Multimedia streaming generally describes a process for allowing access to streamed multimedia content provided by or originating from one or more sources. Increased usage of the Internet has resulted in an increased demand for multimedia streaming and a corresponding need for cost effective multimedia communications solutions.

## SUMMARY

In one general aspect, performance of a network system may be improved by a switch that includes a communications interface to receive a stream of data units that each include a payload portion and an attribute portion, a buffer structured and arranged to store at least payload portions of the data units, a replicator structured and arranged to duplicate at least the payload portion of one or more of the data units, and a second communications interface structured and arranged to enable access to the stream of data units by two or more terminals.

Implementations may include one or more of the following features. For example, the attribute portions of the data units may include IP packet information and/or one or more pieces of layer three information. The switch may include a processor that generates and associates an attribute portion with a payload portion that has been duplicated by the replicator. The replicator may be structured and arranged to duplicate the payload portion of the attribute portion, and the processor may be structured and arranged to change an IP header on one or more of the attribute portions duplicated by the replicator.

The processor may be structured and arranged to change an IP destination address, which may be, for example, the IP address corresponding to the terminal to which access to the payload portion is enabled using a second communications interface. The processor may

be structured and arranged to specify destination information that differs among duplicated payload portions. The data unit may include, for example, audio, video, and streamed media content.

The switch may be structured and arranged to receive a request to receive a stream of data units from a terminal. The switch then may transmit the stream of data units to the requesting terminal. Implementations may include receiving the request from a device other than the terminals that will receive a stream of data units, and the switch may transmit a stream of data units to the requesting terminals or devices. The switch may enable access to the same stream of data units at two different temporal offsets. The switch may include a buffer with more than one pointer to enable access to the stream of data units at two different points and times. The buffer may store more than one instance of the stream of data units.

A source system may interface with the switch. The source system may be capable of duplicating data units and may transmit a stream of data units to the switch. Likewise, one or more terminals may interface with the switch and may receive duplicated data units from the switch.

Implementations may include a system capable of achieving the above features, including, for instance, a source system, a switch, a terminal, and a network between these components. Implementations also may include a sequence of steps performed on the switch, the source system and/or the terminal to achieve these features.

Other features and advantages will be apparent from the following description, including the drawings, and the claims.

## DESCRIPTION OF DRAWINGS

Fig. 1 is a block diagram illustrating a communications system operable to use a switch to generate multiple streams of data units from a single data source.

Fig. 2A is a diagram of an exemplary data unit that may be transmitted in a communications system, such as that shown in Fig. 1.

Fig. 2B is a diagram illustrating one example of the structure of an attribute portion of the data unit of Fig. 2A.

Fig. 3 is a flowchart illustrating a method of transmitting data units in a communications system, such the communications system of Fig. 1.

Fig. 4 is a flowchart illustrating movement of an IP packet through the communications system of Fig. 1.

### DETAILED DESCRIPTION

For illustrative purposes, Figs. 1-4 describe a communications system for implementing techniques for streaming data units. For brevity, several elements in the figures described below are represented as monolithic entities. However, as would be understood by one skilled in the art, these elements each may include numerous interconnected computers and components configured to perform a set of specified operations and/or dedicated to a particular geographical region.

Referring to Fig. 1, communications system 100 may be structured and arranged with a source system 110, two or more terminals 150 and communication software and hardware that enable communications between source system 110 and terminals 150. More particularly, the communications system 100 typically includes the source system 110, a network 120, a switch 130, a network 140 and terminals 150. As will be described in greater detail with respect to Fig. 3, the source system 110 generally transmits one or more data units in a stream of data units across network 120 to one or more switches 130, where data units or portions of data units, are duplicated and transmitted to two or more terminals 150 through network 140.

Typically, a source system 110 may be structured and arranged to convert a media source (e.g., a video or audio feed) into data units for transmission across a network 120. The source system 110 may include a general-purpose computer having a central processor unit (CPU), and memory/storage devices that store data and various programs such as an operating system and one or more application programs. Other examples of a source system 110 include a workstation, a server, a device, a special purpose device or component, a broadcast system, other equipment, or some combination thereof capable of responding to and executing instructions in a defined manner. The source system 110 also typically includes an input/output (I/O) device (e.g., one or more devices having video and audio input and conversion capability), and peripheral equipment such as a display communications card or device (e.g., a modem or a network adapter) for exchanging data with the network 120.

A communications link 115 is used to communicate data between source systems 110 and network 120. Communications link 115 may include, for example, a telephone line, a wireless network link, a cable network, or a direct connection.

The network 120 typically includes hardware and/or software capable of enabling direct or indirect communications between the sending system 110 and the switch 130. The network 120 may include a direct link between the source system 110 and the switch 130, or it may include one or more networks or subnetworks between them (not explicitly shown). Each network or subnetwork may include, for example, a wired or wireless data pathway capable of carrying and receiving data. Examples of network 120 include the Internet, the World Wide Web, a WAN ("Wide Area Network"), a LAN ("Local Area Network"), an analog or a digital wired and wireless telephone network (e.g., a PSTN ("Public Switched Telephone Network"), an ISDN ("Integrated Services Digital Network"), or a xDSL ("any form of Digital Subscriber Loop")), and/or a radio, television, cable, satellite, or any other delivery mechanism for carrying data.

The switch 130 typically is structured and arranged to receive the stream of data units from the source system 110, to duplicate the stream of data units, and to transmit a stream of duplicated data units to one or more terminals 150.

In some implementations, the switch 130 is structured and arranged to perform filtering and forwarding between different domains at the same level of the protocol stack in the OSI ("Open System Interconnection") reference model. For example, in some networks, switch 130 may forward Ethernet frames between different Ethernet segments. In another example, switch 130 may forward IP packets between different IP subnets.

Generally, switch 130 includes a device that performs network operations and functions in hardware (e.g., a chip or part of chip). In some implementations, the device may include an ASIC ("Application Specific Integrated Circuit") implementing network operations logic directly on a chip (e.g., logical gates fabricated on a silicon wafer then manufactured into a chip). For example, an ASIC chip may perform filtering by receiving a packet, examining the IP address of the received packet, and filtering based on the IP address by implementing a logical gate structure in silicon.

Implementations of the device included in the switch 130 may use a Field Programmable Gate Array (FPGA). A FPGA is generally defined as including a chip or chips fabricated to allow third party designers to implement a variety of logical designs (e.g.,

group of gates) on the chip. For example, a third party designer may load a design within a FPGA to replace the received IP addresses with different IP addresses, or may load a design within the FPGA to segment and reassemble IP packets as they are modified while being transmitted through different networks.

Implementations of the device included in the switch 130 may include a network processor. A network processor is generally defined to include a chip or chips for allowing software to specify which network operations will be performed. A network processor may perform a variety of operations. One example of a network processor may include several interconnected RISC ("Reduced Instruction Set Computer") processors fabricated in a network processor chip. The network processor chip may implement software on some of the RISC processors to change an IP address of an IP packet. Other RISC processors in the network processor may implement software that determines which terminals are receiving an IP stream.

Although various examples of network operations were defined with respect to the different devices, each of the devices tends to be programmable and capable of performing the operations of the other devices. For example, the FPGA device is described as the device used to replace IP addresses and segment and reassemble packets; however, a network processor and ASIC are generally capable of performing the same operation.

Data units handled by switch 130 may be accessed by or sent to terminals 150 through network 140. As such, network 140 is structured and arranged to receive data units transmitted from the switch 130 for transmission to the terminals 150.

The network 140 may include hardware and/or software capable of enabling direct or indirect communications between the switch 130 and the terminal 150. As such, the network 140 may include a direct link between the switch 130 and the terminal 150, or it may include one or more networks or subnetworks between them (not shown). Each network or subnetwork may include, for example, a wired or wireless data pathway capable of carrying and receiving data. Examples of the delivery network include the Internet, the World Wide Web, WANs, LANs, analog or digital wired and wireless telephone networks (e.g., PSTN, ISDN, or xDSL), radio, television, cable, satellite, and/or any other delivery mechanism for carrying data. Network 120 and network 140 may share one or more hardware or software devices.

The terminal 150 may include one or more devices capable of receiving the stream of data units transmitted by switch 130 through network 140. The terminal 150 may include a controller (not shown) that processes instructions received from or generated by a software application, a program, a piece of code, a device, a computer, a computer system, or a combination thereof, to direct operations of the terminal 150. The instructions may be embodied permanently or temporarily in any type of machine, component, equipment, storage medium, or propagated signal that is capable of being delivered to the terminal 150 or that may reside with the controller at the terminal 150. The terminal 150 may include a general-purpose computer (e.g., a personal computer) capable of responding to and executing instructions in a defined manner, a workstation 152, a laptop 154, a PDA ("Personal Digital Assistant") 156, a wireless phone, a component, other equipment, or some combination of these items that is capable of responding to and executing instructions.

For instance, in one implementation, the terminal 150 includes one or more information retrieval software applications (e.g., a browser, a mail application, an instant messaging client, an Internet service provider client, or an AOL TV or other integrated client) capable of receiving one or more data units. The information retrieval applications may run on a general purpose operating system and a hardware platform that includes a general purpose processor and specialized hardware for graphics, communications and/or other capabilities. In another implementation, terminal 150 may include a wireless telephone running a micro-browser application on a reduced operating system with general purpose and specialized hardware capable of operating in mobile environments.

Referring to Fig. 2A, an example of a data unit 200A includes an attribute portion 210A and a payload 220A. The attribute portion 210A typically includes parameters that are modified as the data unit 200A moves through the communications system 100. For example, an IP data unit may change the layer two addressing information as the data unit moves through the network (e.g., network 140). The payload 220A typically includes information including content and/or parameters that are intended for communication from the source system and tend to change infrequently during transmission to preserve integrity of the data being transmitted. For example, the payload 220A may include audio or video content being transmitted to a personal computer (e.g., terminal 150).

Referring to Fig. 2B, the attribute portion 210B of the data unit 200B may include several distinct or intermingled fields, such as, for example, fields of an IP ("Internet

Protocol”) packet. More particularly, for example, attribute portion 210B may include a destination address 213B, a source address 215B, a port number 217B indicating the type of IP traffic (e.g., UDP (“User Data Protocol”), TCP (“Transmission Control Protocol”), ICMP (“Internet Control Message Protocol”)), and other parameters at various portions as represented by 211B and 219B. These fields may be arranged as shown or otherwise to accommodate various protocols. The payload in the IP data unit 200B is represented by 220B.

Fig. 3 illustrates a method of transmitting a flow of data units in a communications system 300. Communications system 300 generally includes a source system 310, a switch 320, and terminals 350 and 370. Generally, source system 310 corresponds to the source system 110 in Fig. 1, switch 320 corresponds to the switch 130 in Fig. 1, and terminals 350 and 370 correspond to components of the terminal 150 in Fig. 1.

Initially, source system 310 generates or enables access to a stream of data units in response to a request or otherwise (step 313). When the stream is derived from an analog input (e.g., analog audio or analog video), this may include converting the input into a stream of data units. Even if the input already is in digital form, enabling access to a stream of data units may involve reformatting the input into an appropriately formatted stream of data units. The stream of data units may include a variety of media streams (e.g., video, audio, images, text, and chat). Typically, the data units are IP (Internet Protocol) packets for transmission across the Internet. The source system 310 then transmits the generated stream of data units to one or more switches 320 (e.g., across network 120) (step 316).

When the switch 320 receives the data units from the source system 310 (step 322), it typically buffers and duplicates at least the payload portions of the data units (step 324). In some implementations, data unit attributes may be removed before buffering the payload. For example, for communications over the Internet, switch 320 may remove the IP header information and store only the payload and/or information represented by the payload. As the data unit is being transmitted to terminals, switch 320 modifies the existing header or adds a new IP header to the payload. The switch may use destination information corresponding to the IP address of the destination terminals (e.g., terminals 350 and 370) or an intermediary between switch 320 and the destination terminals as the destination information. With the newly acquired destination information, the switch 320 transmits the

IP packets across the network (e.g., network 140 in Fig. 1) to the appropriate terminal, for example, terminal 370 in Fig. 3 (step 326).

The terminal 370 receives the stream of data units sent by the switch 320 across the network (step 372) and converts the data units into a perceivable output, manipulates the data units, or forwards the data units to another device. For example, a terminal 370 may receive a stream of IP data units containing video content and display or otherwise manipulate the video content on a personal computer or other device.

In some implementations, terminals may be capable of requesting a stream. For example, in Fig. 3, although terminal 370 is shown as automatically receiving a stream of data units (step 372), terminal 350 may be capable of generating a request to receive the stream of data units (step 352). When the switch 320 receives the request (step 328), it may transmit one or more streams of data units to terminal 350 (step 330). Terminal 350 then receives the transmitted streams of data units (step 354).

In some implementations, the switch 320 and the terminals 350 and 360 may be required to communicate periodically to maintain streaming. For example, a terminal 350 may send a "keep stream alive" message every 10 seconds to indicate that the terminal 350 wishes to receive a stream during its transmission.

While receiving a stream (step 354), the terminal 350 may request another stream (step 352) or may request termination of transmission by transmitting a request to stop sending the stream of data units (step 356). Upon receipt of a termination request or upon detection of the conclusion of the stream (step 332), the switch 320 stops transmitting the stream of data units (step 334).

Terminal 350 may elect to resume transmission of the stream of data units and may communicate this election to switch 320 (step 358). When the switch 320 receives a resume instruction (step 336), the switch resumes transmission of the stream (step 338) for receipt by the terminal 350 (step 360).

In some implementations, when the switch 320 receives a stream of IP data units from source system 310, the switch 320 may duplicate the payload portion of the data unit, and may modify or replace the IP destination address of the IP data units to specify the IP addresses of one or more conference terminals.

Typically, the switch 320 may duplicate any particular payload or all payloads more than once such that the payload of a received data unit is transmitted as the payload in several

or all data units generated by the switch 320. Moreover, the switch 320 may receive a data unit and transmit the payload as two or more data units.

Conversely, the switch 320 may receive and combine more than one data unit and transmit the combined payloads of the multiple data units in a single data unit, or the switch 320 may divide the payload of a received data unit and combine the divided portions with the payloads of two or more other received data units.

As will be described with respect to Fig. 4, since the payloads are buffered, the switch 320 may produce several instantiations of the buffered payload, simultaneously, with or without offset, and may transmit those different instantiations to the same or different terminals as needed. These instantiations may be streamed to different time offsets so that terminal 350 may be listening to a stream of data units at a different point than terminal 370. For example, terminal 350 may be receiving the same song on an Internet radio station as terminal 370, but terminal 350 may be receiving data units that correspond to the song thirty (30) seconds later in time.

Fig. 4 illustrates one example of the flow of IP data units from a source system 410 to terminals 470, 480 and 490. The source system 410 with an IP address of X.1.1.1 transmits an IP packet 420 to the switch 430 with a destination address of Y.1.1.1. The IP packet 420 includes the switch address as a destination address 424 and the source system address as a source address 426. The IP packet 420 also may include other header information 422 and a payload 428.

The switch 430, with a source address of Y.1.1.1, duplicates at least the payload 428 of the IP packet 420 and transmits an IP data unit 440 to terminal 470 with an address of Z.1.1.1. The IP data unit 440 includes the source IP address 446 of the switch 430 with a destination address 444 of terminal 470. The IP data unit 440 may include additional header information 442 and a payload 448. Similar activities take place to generate IP data units 450 and 460 for terminals 480 and 490.

Other implementations are within the scope of the following claims. For instance, implementations may include a manager requesting transmission of the stream to the terminals. The manager may include a network operator, a managing server, a workstation, or a scheduling agent. For example, a managing workstation may request transmission of the stream of data units to terminals beginning at a certain time.

Implementations also may include storing more than one instance in the stream of data units in the buffer. For example, the switch may store more than one copy of the stream to accommodate large numbers of requests. In another example, the stream may buffer the same stream of data units at two different spots.

Implementations also may include having the switch maintain state information on one or more requesting users. For example, the switch may establish or assist other devices in establishing a profile for the terminal that receives the stream of data units. In this manner, the switch leverages information available to network processors and devices, which may not ordinarily be accessible to servers, in establishing profiles.

Still further, implementations may include monitoring switch access and usage levels. For example, the switch may determine that the switch is operating at 85% of maximum performance in a given metric (e.g., processor use, memory use, number of users, bandwidth).

Implementations may include translating or transcoding content between different formats. For example, the switch may adjust transmission of packets to minimize the bandwidth consumed. In another example, the switch may convert a proprietary video-encoding format to a standards-based encoding format.

Implementations also may include inserting one or more pieces of content in an existing stream of data units. For example, the switch may insert video stream advertisements in a transmission of other video content. The inserted video content may reside on the switch, or it may be accessed from another device distinct from the switch that is transmitting the stream of data units. Implementations of inserting content may include creating tags associate with certain pieces of content being replicated. For example, the switch may be replicating a television show with commercials tagged in the video content periodically. Upon receiving the tag, the switch may stop transmitting the television show and access a series of advertisements. Implementations may include having the tag designate which advertisement to access. For example, the tag may indicate the device on which the advertisement is located. Implementations of tags also may enable accessing more than one selection of advertisements. The selection of advertisements may depend on one or more variables including, but not limited to, terminal information and profiling, the state of the network and other factors.

In addition, implementations may include inserting content based on information maintained about the user profile and user state. For example, the switch may determine that a terminal is at a suitable point in a content piece to receive an inserted advertisement for users with a particular profile.